

A New Dynamical Core with Local Mesh Refinement

Todd Ringler

Los Alamos National Laboratory
Theoretical Division

LA-UR 09-05207

Climate, Ocean and Sea-Ice Modeling Project
<http://public.lanl.gov/ringler/ringler.html>

Collaborators and Contributors:

Max Gunzburger

Bill Skamarock

Michael Duda

John Thuburn

Joe Klemp

Lili Ju



Background

Two lines of research have come together to create this dynamical core. This effort started about two years ago.

1. Creating a C-grid staggering that is applicable to a wide class of meshes that conserves energy, conserves potential vorticity and dissipates potential enstrophy with the appropriate time scale. Applicable (without modification) to variable resolution meshes.

- Ringler, T., J. Thuburn, J. Klemp and W. Skamarock, 2009: A unified approach to energy conservation and potential vorticity dynamics on arbitrarily structured C-grids, Journal of Computational Physics, accepted. ([pdf](#)).
- Thuburn, J., T. Ringler, J. Klemp and W. Skamarock, 2009: Numerical representation of geostrophic modes on arbitrarily structured C-grids, Journal of Computational Physics, 2009: 228 (22), 8321-8335. doi:10.1016/j.jcp.2009.08.006 ([link](#))

2. The creation of high quality variable resolution meshes that can be used with the above numerical method.

- Ju, L., T. Ringler and M. Gunzburger, 2009, Voronoi Diagrams and their Application in Climate Science, Numerical Techniques for Global Atmospheric Models, Lecture Notes in Computational Science, draft. ([pdf](#)).
- Ringler, T., L. Ju and M. Gunzburger, 2008, A multiresolution method for climate system modeling: application of spherical centroidal Voronoi tessellations, Ocean Dynamics, 58 (5-6), 475-498. doi:10.1007/s10236-008-0157-2 ([link](#))

Background (continued)

A one-year infusion of funds from DOE Office of Science has allowed us to make a strong push to move this from a set of good ideas to a prototype ocean (and atmosphere) dynamical core.

Examples include:

1. Implementing high-order transport (Lowerie and Buoni)
2. Sub-grid closures for variable resolution meshes (Carlson)
3. Full implementation of immersed boundary method (Mohd-Yosuf)
4. Dynamical core formulation and implementation (Ringler)
5. Implementation of method as a option for the CCSM CAM (coupling to CCSM led by Mirin et. al at LLNL, model formulation led by Skamarock et. al at NCAR MMM. NCAR MMM is an un-funded partner for this one-year effort.)

Cutting to the chase on the numerical method

Analytic results for the nonlinear shallow-water equations:

(includes full ocean models based on isopycnal coordinates)

1. Stationary geostrophic mode is recovered.
2. Total energy is conserved to within time truncation.
 - a. Coriolis force is energetically-neutral
 - b. Transport of KE is conservative
 - c. KE/PE exchange is equal and opposite.
3. Potential vorticity is conserved to round-off. PV is compatible with an underlying thickness equation.
4. Potential enstrophy can be dissipated through the discretization of potential vorticity transport (i.e. monotone transport).

Results hold for a wide class of meshes: Lat/Lon, Stretched Lat/Lon, Voronoi Tessellations, Delaunay Triangulation and Conformally-mapped cubed sphere meshes.

An important aspect of this method is the control over PV dynamics that we can exert while discretizing the momentum equation.

$$\mathbf{k} \cdot \nabla \times \left[\frac{\partial \mathbf{u}}{\partial t} + q(h\mathbf{u}^\perp) = -g\nabla(h + h_s) - \nabla K \right]$$

nonlinear Coriolis force

$$\frac{\partial \eta}{\partial t} + \mathbf{k} \cdot \nabla \times [\eta \mathbf{u}^\perp] = 0$$

$$\frac{\partial \eta}{\partial t} + \nabla \cdot [\eta \mathbf{u}] = 0$$

potential vorticity flux

$$\frac{\partial(hq)}{\partial t} + \nabla \cdot [hq\mathbf{u}] = 0$$

The nonlinear Coriolis force IS the PV flux in the direction perpendicular to the velocity. So the evolution of PV is controlled entirely through the discretization of the nonlinear Coriolis force.

This insight was also helpful in the formulation of an extension of GM to include the influence of the Bolus velocity on PV transport.


• Ringler, T. and P. Gent, 2009: An eddy closure for potential vorticity, Journal of Physical Oceanography, submitted. ([pdf](#)).

Defining the discrete system

$$\frac{\partial h_i}{\partial t} + \left[\nabla \cdot \left(\hat{h}_e u_e \right) \right]_i = 0$$

$$\frac{\partial u_e}{\partial t} + \hat{q}_e [hu]_e^\perp = \left[\nabla (gh_i + K_i) \right]_e$$

- locations of velocity points
- ▲ locations of potential vorticity
- locations of thickness points




As with all C-grid methods, the biggest challenge is how to reconstruct the nonlinear Coriolis force.

- 1) How to compute $[hu]_e^\perp$?
- 2) How to compute \hat{q} ?

Reconstructing the nonlinear Coriolis force

(recall that the nonlinear Coriolis force is the the PV-flux perpendicular to the velocity)



$$\frac{\partial u_e}{\partial t} + \hat{q}_e [hu]_e^\perp = [\nabla (gh_i + K_i)]_e$$

$$\frac{\partial u_e}{\partial t} + Q_e^\perp = [\nabla (gh_i + K_i)]_e$$

$$d_e Q_e^\perp = \sum_j w_e^j l_j F_j \bar{q}_{ej}$$

$$F_j = \hat{h}_j u_j \quad \text{thickness flux}$$


$$w_e^j = -w_j^e \quad \text{weights are equal and opposite}$$

$$\bar{q}_{ej} = \bar{q}_{je} \quad \text{PV is symmetric}$$

The nonlinear Coriolis force will be energetically neutral for any \bar{q}_{ej} .
This is an extension to what Sadourny (1975) showed for regular meshes.

Reconstructing the nonlinear Coriolis force

(recall that the nonlinear Coriolis force is the the PV-flux perpendicular to the velocity)



$$\frac{\partial u_e}{\partial t} + \hat{q}_e [hu]_e^\perp = [\nabla (gh_i + K_i)]_e$$

$$\frac{\partial u_e}{\partial t} + Q_e^\perp = [\nabla (gh_i + K_i)]_e$$

Energy conservation is obtained by:

$$\bar{q}_{ej} = \bar{q}_{je} = \frac{1}{2} (\hat{q}_e + \hat{q}_j)$$

Energy is conserved for any \hat{q}_e and any \hat{q}_j !

The APVM is obtained by:


$$\hat{q}_e = \frac{1}{2} (q_{v1} + q_{v2}) - \underbrace{\frac{dt}{2} \mathbf{u} \cdot \nabla q}_{\text{upstream bias}}$$

This is the upstream bias that leads to dissipation of potential enstrophy.

(In terms of computational cost, this closure is essentially free.)


SWTC#5, 100 day integration, quasi-uniform mesh

potential enstrophy dissipation with 100 day time scale
energy conserved to $1e-8$



<http://public.lanl.gov/ringler/movies/2009/APVM.mov>


Globally-average potential enstrophy when using the APVM



Variable-resolution simulations


We would like to have the ability to statically refine the mesh in order to improve the local (and global?) simulation.

This improvement arises primarily from the regional improvement of specific physical processes (e.g. cloud-process in the atmosphere and geostrophic-eddy processes in the ocean.)




Variable Resolution Mesh

(oceanographers: please shift mesh ~120 degrees to the east!)




Variable Resolution Mesh

(oceanographers: please shift mesh ~120 degrees to the east!)




Variable Resolution Mesh

(oceanographers: please shift mesh ~120 degrees to the east!)




Variable Resolution Mesh

(oceanographers: please shift mesh ~120 degrees to the east!)



Variable Resolution Mesh

(oceanographers: please shift mesh ~120 degrees to the east!)




SWTC#5, 100 day integration, variable-res mesh

potential enstrophy dissipation with 100 day time scale
energy conserved to $1e-8$

NOTE: The grid-scale phenomena at $t=0$ is due to geostrophic adjustment, i.e. the numerical model responding to imperfect balance at $t=0$. As appropriate, this results in gravity and Rossby waves that disperse as the flow attains geostrophic balance.

NOTE: Each grid cell is drawn as polygon and shaded with a single color. Thus, data is shown without interpolation. No dissipation or other “closure” is used, including in the mesh transition zone.



http://public.lanl.gov/ringler/movies/2009/APVM_var.mov

Double-Gyre Problem

5 km mesh



5000 km x 2500 km domain

Forcing consistent with Greatbatch and Nadiga (2000).

No bottom drag

Laplacian mixing on velocity at $1.0 \text{ m}^2/\text{s}$
no-slip boundary conditions

NOTE: The purpose here is to confirm that the numerical scheme can handle strong eddy activity without the need for stabilization through ad hoc dissipation. The scheme is stable for a very wide range of dissipations, even those that clearly lead to an underdamped system.



Potential Vorticity: one frame per day

http://public.lanl.gov/ringler/movies/2009/eddies_18.mov


How is this implemented into a
high-performance computing environment?

Decomposition for implementation on MPPs.

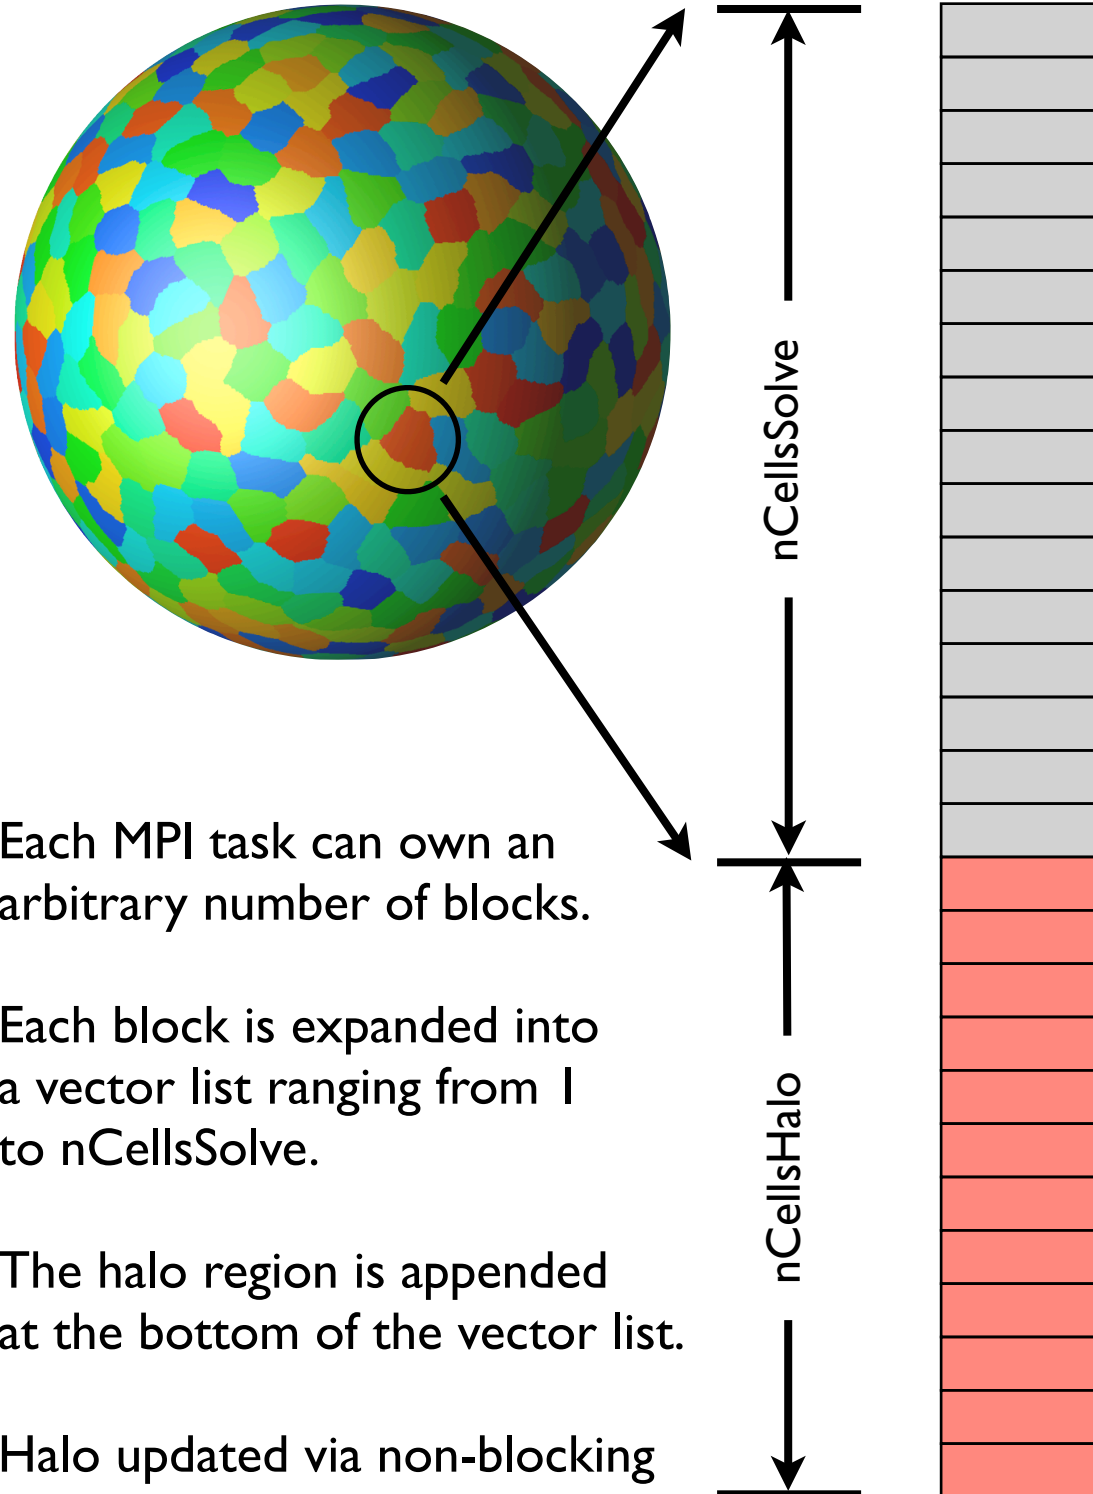
Horizontal decomposition is done with METIS.

At present, each processor gets one block. In the future, each processor will own multiple blocks to improve cache reuse.

We will be supporting multiple climate system model components under this single software framework: atmosphere, ocean, and (in all likelihood) ice sheet.



Data Layout: Unstructured in horizontal, structured in vertical: Creating a data structure to exploit hybrid computing architectures.



$$nCellsTotal = nCellsSolve + nCellsHalo$$

The ordering of the cells is determined by Reverse Cuthill-McKee (RCM) to maximize cache reuse.

The data access pattern is repeated in the vertical leading to arrays dimensioned as

$mass(nVertLevels, nCellsTotal, nBlocks)$

Example: sum mass at neighbors cells, each proc executes

```
do iBlock = 1, nBlocks
  do i = 1, nCellsSolve(iBlock)
    do j = 1, nCellsOnCell(i, iBlock)
      coc = cellsOnCell(i, j, iBlock)
      do k = 1, nVertLevels
        r(k, i) = r(k, i) + mass(k, coc, iBlock)
      enddo
    enddo
  enddo
enddo
```

The vertical loop is always the inner most loop to vectorize and to limit indirect addressing.

The layout is repeated for variables that live at vertices and at edges.


Preliminary scaling results of SWM with global 30 km mesh.

Super-linear scaling out to 512.

This is about 1000 cells per processor. Yet we have not implemented any flop-intensive parts of the model, such as model physics, high-order transport, or 10 to 100 tracer constituents.

Weak scaling at the level observed with POP (200 horizontal cells per processor) seems readily attainable here. Especially since we anticipate no global reductions.


NOTE: Scaling results reproduced up to 1024 processors on Ranger with global 15 km mesh. (Duda)




Lobo: AMD Opteron cores for computation, using an Infiniband interconnect. 272x16-core nodes for production capacity computing. Each node has 32 GB of non-uniform-access memory. CPUs run at 2.2 Ghz and have 0.5-MB L2.

Some results from MPI visit - hydrostatic MPAS core


APVM



CONTOUR FROM -0.00005 TO 0.00008 BY 0.00001



Centered PV



CONTOUR FROM -0.00005 TO 0.00008 BY 0.00001




Jablonski and Williamson baroclinic-wave test case
Relative vorticity (1/s), day 9, at ~ 850 hPa, 10262 cells

MPI, 10 November 2009


From Skamarock, implementation of numerical scheme as a hydrostatic atmospheric dynamical core.

3-D Supercell Simulation ~500 m Horizontal Grid

Rectangular Grid



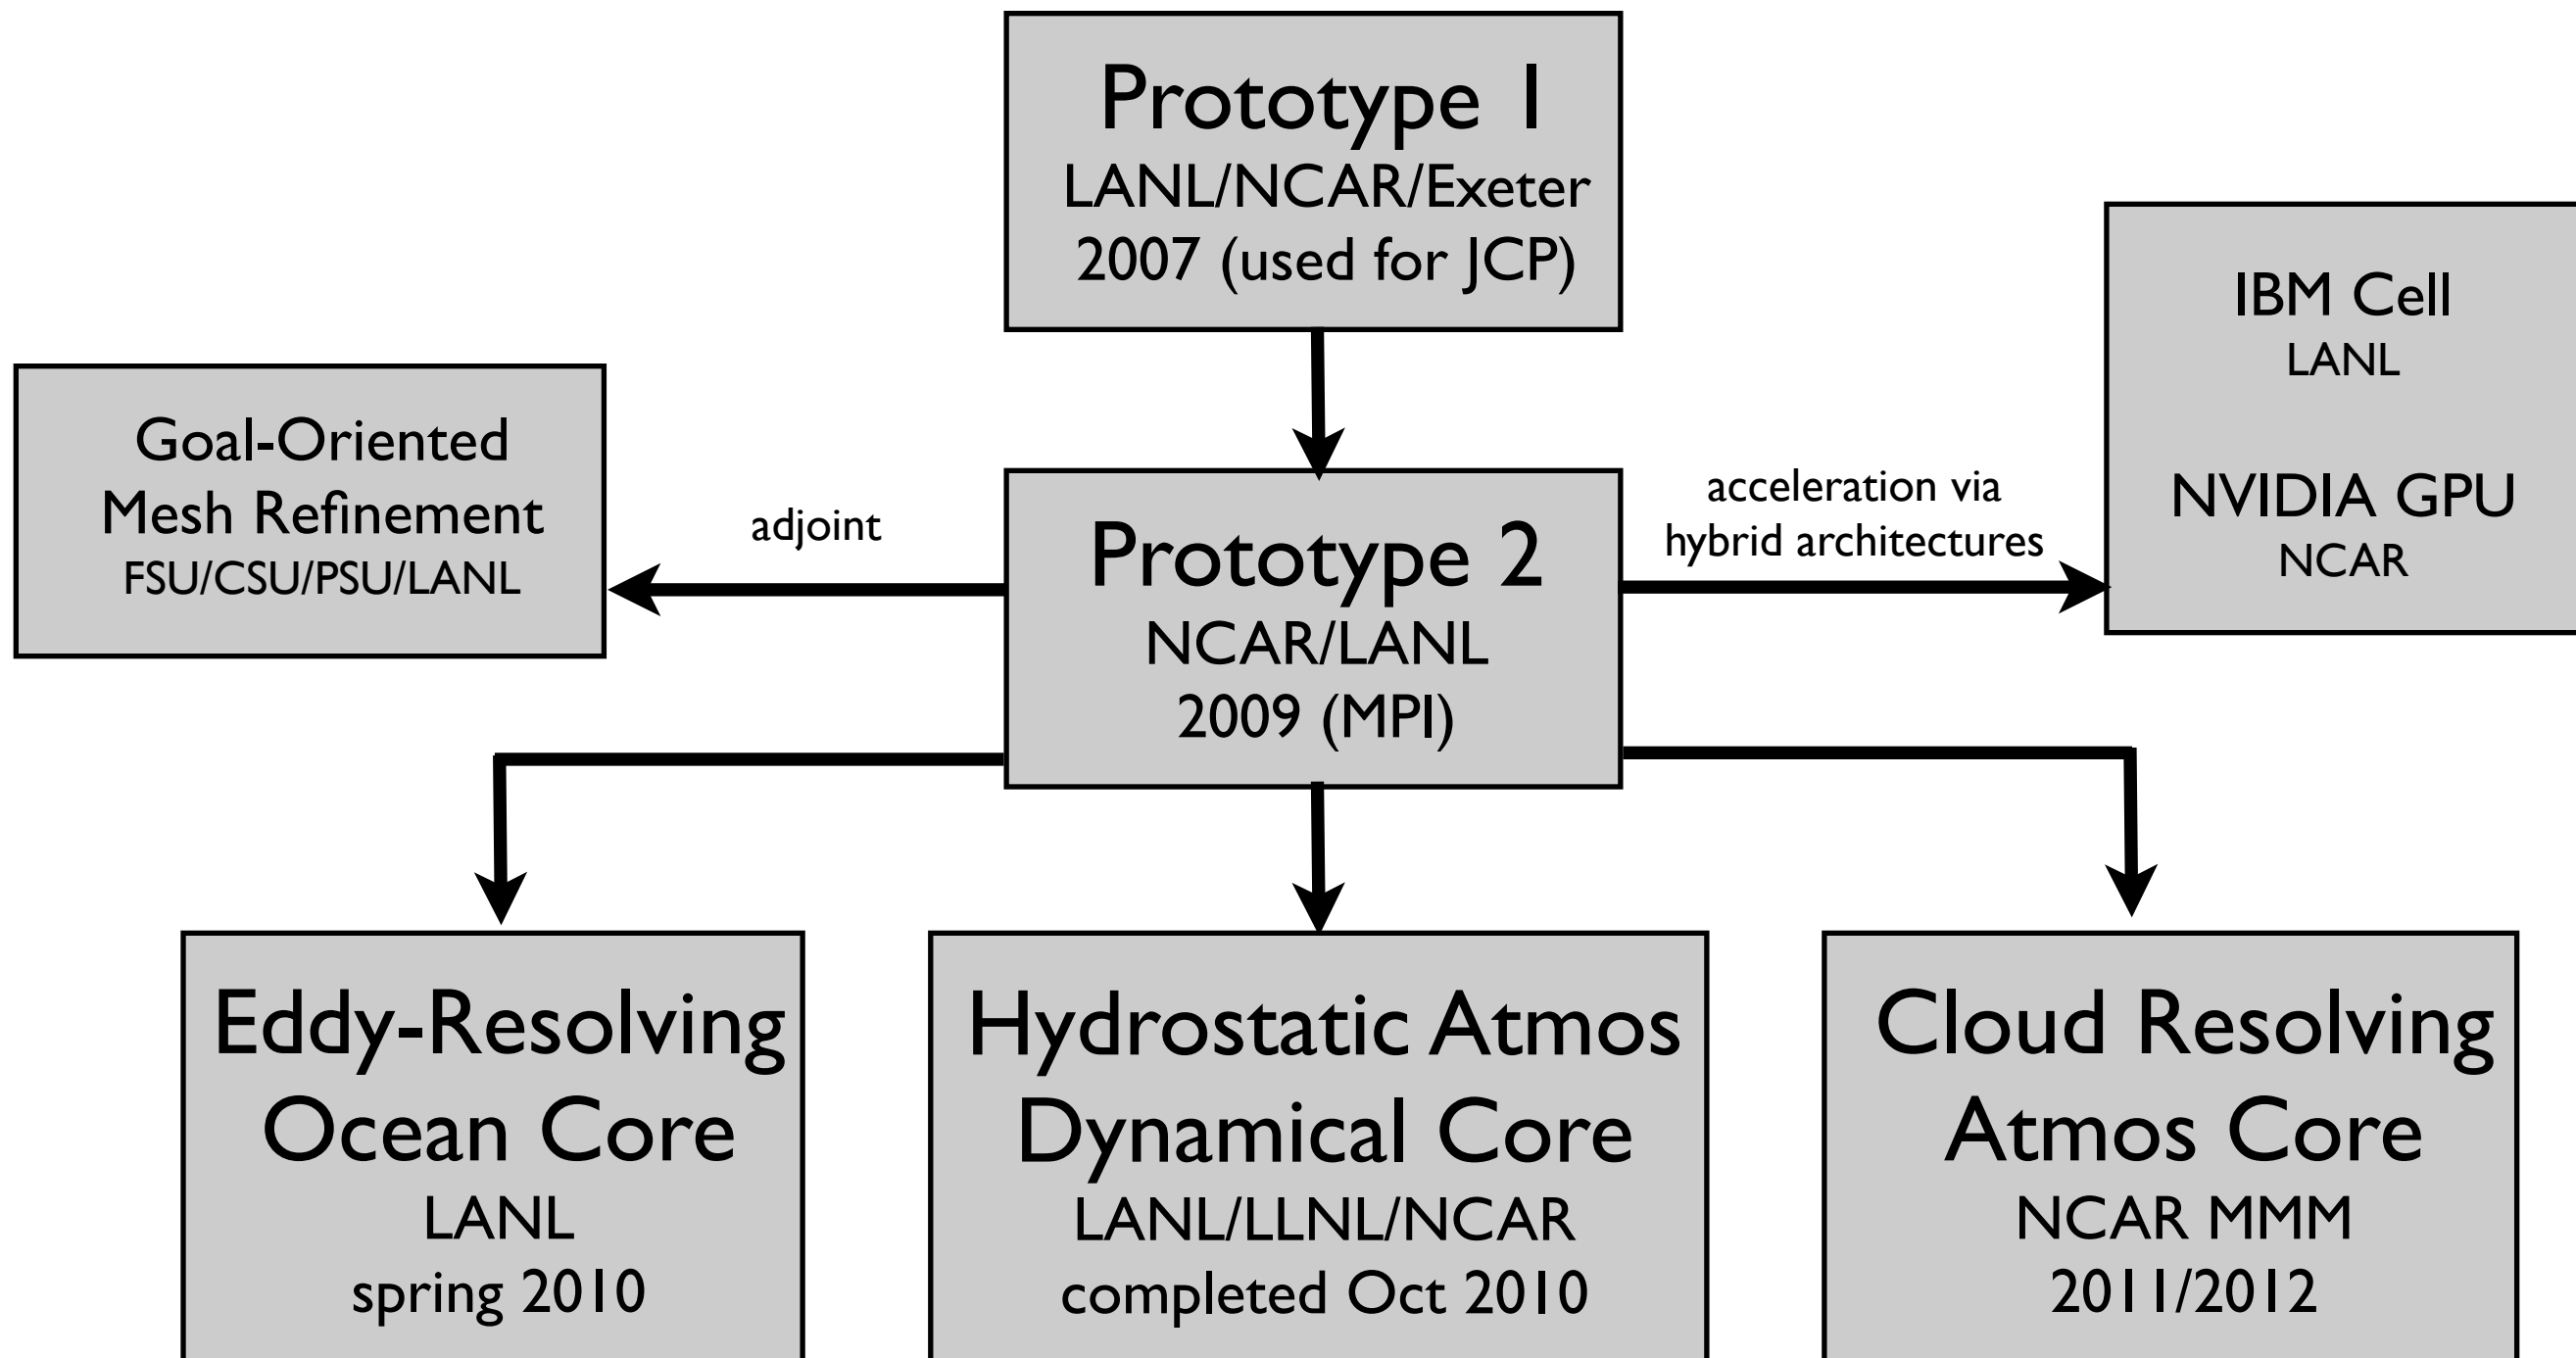
Hexagonal Grid



- Vertical velocity contours at 1, 5, and 10 km (c.i. = 3 m/s)
- 30 m/s vertical velocity surface shaded in red
- Rainwater surfaces shaded as transparent shells
- Perturbation surface temperature shaded on baseplane

From Klemp, implementation of numerical scheme as a cloud resolving model

Proposed Development Path



We are proposing is a next-generation, ocean dynamical core that is backward-compatible to POP/HYPOP.

We are proposing is a next-generation, ocean dynamical core that is backward-compatible to POP/HYPOP.

What is the method?

1. The method is a C-grid staggering that conserves total energy and PV. In addition it allows for the direct control over the evolution of PV through the discretization of the nonlinear Coriolis force.
2. The method is applicable to a very wide class of meshes, including variable resolutions meshes.

We are proposing is a next-generation, ocean dynamical core that is backward-compatible to POP/HYPOP.

What is the method?

1. The method is a C-grid staggering that conserves total energy and PV. In addition it allows for the direct control over the evolution of PV through the discretization of the nonlinear Coriolis force.
2. The method is applicable to a very wide class of meshes, including variable resolutions meshes.

What are the benefits?

1. The method can be implemented on POP stretched grid with no modification, so backward compatibility in terms of meshes, boundary conditions, etc. is supported.
2. The ability to conduct regionally eddy-resolving simulations within the framework of a global model.
3. The ability to more precisely position degrees of freedom to improve simulation, for example equatorial regions and overflow regions.
4. The opportunity to incorporate (either explicitly or conceptually) a set of forward-looking ideas related to hybrid computing, LANS-alpha sub-grid closure, high-order transport, immersed boundary method, layered modeling and implicit time stepping.
5. Tighter integration (at the software foundation) with other climate model components, hence the opportunity to quickly leverage progress made by others.

We are proposing is a next-generation, ocean dynamical core that is backward-compatible to POP/HYPOP.

What is the method?

1. The method is a C-grid staggering that conserves total energy and PV. In addition it allows for the direct control over the evolution of PV through the discretization of the nonlinear Coriolis force.
2. The method is applicable to a very wide class of meshes, including variable resolutions meshes.

What are the benefits?

1. The method can be implemented on POP stretched grid with no modification, so backward compatibility in terms of meshes, boundary conditions, etc. is supported.
2. The ability to conduct regionally eddy-resolving simulations within the framework of a global model.
3. The ability to more precisely position degrees of freedom to improve simulation, for example equatorial regions and overflow regions.
4. The opportunity to incorporate (either explicitly or conceptually) a set of forward-looking ideas related to hybrid computing, LANS-alpha sub-grid closure, high-order transport, immersed boundary method, layered modeling and implicit time stepping.
5. Tighter integration (at the software foundation) with other climate model components, hence the opportunity to quickly leverage progress made by others.

What are the disadvantages?

1. Getting use to a new software framework.

Manuscripts available ...

<http://public.lanl.gov/ringler/publications.html>

- Ringler, T., J. Thuburn, J. Klemp and W. Skamarock, 2009: Numerical treatment of energy and potential vorticity on arbitrarily structured C-grids, Journal of Computational Physics, accepted. ([pdf](#)).
- Thuburn, J., T. Ringler, J. Klemp and W. Skamarock, 2009: Numerical representation of geostrophic modes on arbitrarily structured C-grids, Journal of Computational Physics, 2009: 228 (22), 8321-8335. doi:10.1016/j.jcp.2009.08.006 ([link](#))
- Ju, L., T. Ringler and M. Gunzburger, 2009, Voronoi Diagrams and their Application in Climate Science, Numerical Techniques for Global Atmospheric Models, Lecture Notes in Computational Science, draft. ([pdf](#)).
- Ringler, T., L. Ju and M. Gunzburger, 2008, A multiresolution method for climate system modeling: application of spherical centroidal Voronoi tessellations, Ocean Dynamics, 58 (5-6), 475-498. doi: 10.1007/s10236-008-0157-2 ([link](#))
- Ringler, T. and P. Gent, 2009: An eddy closure for potential vorticity, Journal of Physical Oceanography, submitted. ([pdf](#)).

extra

Summary (focused on numerical method).

Summary (focused on numerical method).

I. The anticipated potential vorticity method allows for the realistic simulation of PV-dynamics while conserving total energy. (We will want to dissipate energy, but now we can choose the appropriate closure based on physical arguments, not based on the need to stabilize the numerical scheme.)

Summary (focused on numerical method).

1. The anticipated potential vorticity method allows for the realistic simulation of PV-dynamics while conserving total energy. (We will want to dissipate energy, but now we can choose the appropriate closure based on physical arguments, not based on the need to stabilize the numerical scheme.)
2. We have demonstrated a finite-volume scheme that allows for the use of variable resolution without the need for ad hoc regularization at the nest and/or transition zone. (We will want to modify the numerics in the transition zone, but now we have the luxury of addressing the closure without having to consider numerical stability.)

Summary (focused on numerical method).

1. The anticipated potential vorticity method allows for the realistic simulation of PV-dynamics while conserving total energy. (We will want to dissipate energy, but now we can choose the appropriate closure based on physical arguments, not based on the need to stabilize the numerical scheme.)
2. We have demonstrated a finite-volume scheme that allows for the use of variable resolution without the need for ad hoc regularization at the nest and/or transition zone. (We will want to modify the numerics in the transition zone, but now we have the luxury of addressing the closure without having to consider numerical stability.)
3. We expect that this model, when endowed with high-order, DG-like transport, 10s to 100s of tracers and model physics, will scale competitively when compared to other methods. (We expect that the FV-cores and DG-cores will look more and more similar going forward.)

Summary (focused on numerical method).

1. The anticipated potential vorticity method allows for the realistic simulation of PV-dynamics while conserving total energy. (We will want to dissipate energy, but now we can choose the appropriate closure based on physical arguments, not based on the need to stabilize the numerical scheme.)
2. We have demonstrated a finite-volume scheme that allows for the use of variable resolution without the need for ad hoc regularization at the nest and/or transition zone. (We will want to modify the numerics in the transition zone, but now we have the luxury of addressing the closure without having to consider numerical stability.)
3. We expect that this model, when endowed with high-order, DG-like transport, 10s to 100s of tracers and model physics, will scale competitively when compared to other methods. (We expect that the FV-cores and DG-cores will look more and more similar going forward.)
4. The same core numerical algorithm appears appropriate for both eddy-resolving ocean simulations and cloud-resolving atmosphere simulations. (As it should be).